Using intrinsic complexity of turn-taking games to predict people's reasoning strategies
(Bachelorproject)

Gerben Bergwerff, S2077280, gbergwerff@xs4all.nl,
Rineke Verbrugge*, Ben Meijering*, Stefan Wierda*

January 30, 2014

Abstract

We studied the cognitive strategy participants use while playing the Marble Drop Game. The Marble Drop Game is a turn-based game designed to investigate higher-order social reasoning. We analyzed different strategies by fitting models containing different strategies that participants could use for solving the games. We compare all models with actual reaction times of participants using goodness of fit. This results in one model that best explains the reaction times of the participants, and thus is a good explanation for the strategy participants use. We found that the model of a combination of forward reasoning and backward induction is the best predictor for the complexity of the games for participants and thus for the reasoning strategy that participants use.

1 Introduction

Theory of mind (ToM) is what makes us able to reason about other people. We use it to understand a sentence like 'I believe Ann knows that Peter thinks ...' [Verbrugge (2009)] and forms of social interaction, for example playing games. There has been a lot of research about the age when ToM develops [Flobbe, Verbrugge, Hendriks, and Krämer (2008)] [Perner and Wimmer (1985)] and the performance of humans versus computer opponents on ToM tasks [McKelvey and Palfrey (1992)] [Meijering, Van Rijn, Taatgen, and Verbrugge (2011)]. However, little research on ToM concentrates on the cognitive basis for ToM [Apperly (2010)]. Research in cognitive neuroscience shows that ToM reasoning employs many brain regions [Gallagher and Frith (2003)]. Therefore, it is probable that ToM reasoning consists of multiple serial and concurrent cognitive processes. Cost-benefit trade-offs have a cascading effect on cognitive load [Borst, Taatgen, and van Rijn (2010)] and thus also affect ToM reasoning. The strategy people use is shown to affect the cost-benefit trade-offs between cognitive resources [Gray, Sims, Fu, and Schoelles (2006)]. The investigation of strategies is therefore likely to teach us more about the cognitive bases of ToM reasoning [Ghosh and Meijering (2011)].

In a recent study [Meijering, van Rijn, Taatgen, and Verbrugge (2012)], the researchers designed and used a two-player game (the Marble Drop Game) to investigate the ongoing process of ToM reasoning. The current study elaborates on the research done by Meijering et al. by analyzing the research data collected with the Marble Drop Game.

1.1 The Marble Drop Game

The Marble Drop Game is a strategic two-player game, where both players' goal is to earn as many points as possible. In the original Marble Drop Game, as used by Meijering et al., four hues of blue and four hues of orange indicated the payoff. A darker hue in-

*University of Groningen, Institute of Artificial Intelligence

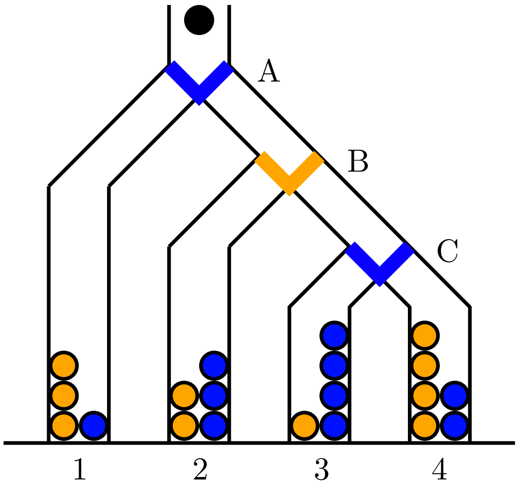

Figure 1: A trial of the Marble Drop Game containing second order reasoning. The blue player has to reason about what side of trapdoor C the orange opponent thinks blue will open. The number of marbles in the bins of the game are isomorphic to the hue used in the original experiment by Meijering et al.

dicated a higher payoff. The hue is isomorphic to a number of points $\in \{1, 2, 3, 4\}$ [Meijering et al. (2012)]. For ease of reading and explaining we will from now on indicate the orange and blue hues as number of orange and blue marbles, where each marble corresponds to one point for the player of that color. In the game, each player controls one or more trapdoors leading to bins containing marbles. Every bin contains one to four marbles of the colors blue and orange, corresponding to points for the blue player and the orange opponent, respectively. A typical Marble Drop Game trial containing second-order reasoning has four bins and three trapdoors (see Figure 1). In this particular game, the blue player can find his highest number of marbles in bin 3 and the orange opponent can find his highest number of marbles in bin 4. Backward induction will always yield the most optimal solution to solve this problem [Heifetz (2012)]. A player using backward induction will start at the bottom of the tree, working his way up to the top. For the structure of the Marble Drop Game, backward induction will always take 6 steps [Szymanik, Meijering, and Verbrugge (2013)]. As an example we apply backward induction to the Marble Drop Game trial in Figure 1: We start at the bottom where the blue player first decides what side of trapdoor C he will choose. When comparing bin 3 with bin 4, blue will choose bin 3 as it contains more blue marbles. Next, blue will decide what the rational orange opponent chooses at trapdoor B: bin 2 or letting the marble get to trapdoor C where the blue player will choose bin 3. Comparing bin 3 with bin 2, bin 2 contains the most marbles for the orange opponent, so the rational orange opponent chooses bin 2 at trapdoor B. Finally, blue will decide what side to choose at trapdoor A, bin 1 or letting the marble get to trapdoor B where orange chooses bin 2. Comparing bin 1 with bin 2, bin 2 contains the most marbles for blue, so the blue player has to choose the side leading to bin 1 at trapdoor A in order to receive the highest possible payoff.

However, eye-tracking data from [Meijering et al. (2012)] showed that participants are more likely to use another form of reasoning called forward reasoning plus backtracking. We will analyze different strategies to predict the complexity of the games.

1.2 Forward Reasoning + Backtracking

Forward reasoning plus backtracking is a combination of forward reasoning [Meijering et al. (2012)] and backward induction. A player using forward reasoning plus backtracking, starts using forward reasoning to find which side of the trapdoors should be opened to get to the bin with the highest score and then uses backward induction to reason if the bin is attainable. In other words: is it possible to get the marble to the bin with the highest payoff or will the rational opponent be more likely to choose an other side of the trapdoor? This leads to the following hypothesis:

Hypothesis 1 - H_{frb}:
The reaction times on games have a significant correlation with the number of steps when solving the game using forward reasoning plus backtracking, thus suggesting the forward reasoning plus backtracking strategy is used.

1.3 Λ-difficulty

A paper about the complexity of the Marble Drop Game [Szymanik et al. (2013)] suggests that the difference in reaction time is mainly due to the structural difficulty of the game tree. They suggest that if a tree has a maximum payoff that is not at the bottom leaves of the tree, the tree can be substituted by a smaller tree, leaving out the leaves after the maximum payoff. To compute the difficulty of a tree Λ^i_{k+1} we look from the perspective of a player $i \in \{1, 2\}$, and the number of k-alterations ($k \leq 0$) between players in the tree, starting at the first node controlled by player i. This is defined in Definition 1:

Definition 1 *Let us assume that the players $\{1, 2\}$ strictly alternate in the game; Let player $i \in \{1, 2\}$.*
Then:

- *In a Λ^i_1 tree, all the nodes are controlled by Player i.*
- *A Λ^i_{k+1} tree, a tree of k-alternations for some $k \geq 0$, starts with a Player i node.*

To compute the minimal subtree that contains the highest possible payoff we use Definition 2 and Definition 3. The Lambda difficulty of this minimal sub-tree using Definition 1 is henceforward called its lambda-degree.

Definition 2 *A game T is generic, if for each player, distinct end nodes have different payoffs.*

Definition 3 *Suppose $i \in \{1, 2\}$. If T is a generic game tree with the root node controlled by Player i and n is the highest possible payoff for Player i, then T^- is the minimal subtree of T containing the root node and the node with payoff n for Player i.*

Trees that have a small minimal subtree should take less time to be solved than trees that have a large minimal subtree. If the highest number of marbles is in one of the top leaves, it should be relatively easy for the player to see if it is possible to get the marble to the bin with the highest payoff. Therefore, reaction times on small minimal subtrees should be lower than reaction times on large minimal subtrees. The trees used in the Marble Drop Game are all of the type Λ^i_2 or Λ^i_3, where the player $i = 1$ and the opponent $i = 2$.

This leads to the following hypothesis:

Hypothesis 2 - H_Λ:
The games having a difficulty of Λ^1_2 have a significantly shorter reaction time than games having difficulty Λ^1_3.

1.4 Research questions

Based on the two hypotheses discussed above, the two questions that will be answered in this study are:

1. Does the H_Λ explain the difference in reaction time between trials of the marble drop game?
2. Does the H_{FRB} explain the difference in reaction time between trials of the marble drop game?

To answer these questions we will re-evaluate the data collected by the research of Meijering et al. (2012). Using statistical methods, we will test whether either of the above hypotheses are valid and thereby answer our research questions. The results of this research will give us more insight about how and why humans choose a reasoning strategy and the cognitive aspect of ToM.

2 Methods

2.1 Participants

Twenty-three first-year psychology students (14 female) with a mean age of 20.8 years (ranging from 18 to 24) participated in the experiment in exchange for course credit. All participants had normal or corrected-to-normal visual acuity. None of the participants had difficulties distinguishing between the two colors (blue and orange) presented in the experiment [Meijering et al. (2012)].

2.2 Experimental design

The participants were asked to solve 20 training trials of increasing difficulty and 64 experimental trials of the marble drop game. The experimental trials were divided over two blocks of 32 trials each. In the first block, 10 of the participants were prompted by asking what side of the trapdoor the opponent would choose. In the second block, none of the participants were prompted. All trials used only items with a payoff structures that required the participants to use second-order reasoning. In total, 16 Different game trees were used as basis for the items. These game trees are henceforward called types. One of these types was chosen randomly for each trial. The payoff structure of all 16 types of the Marble Drop Game can be found in Attachment A.

2.3 Attainability

The highest score (i.e. 4) is not always attainable by the player in every trial. In some trials, the highest value is not attainable as the rational opponent would never choose the side of the trapdoor that will lead to the highest value for the player. For every trial, we decided whether or not the highest value was attainable for the player and we divided the data into two classes: attainable, not attainable, accordingly.

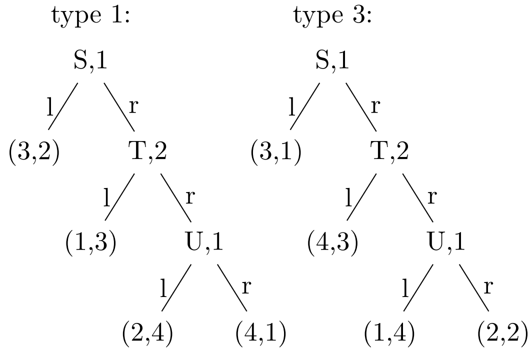

Figure 2: Pay-off structures of type 1 and type 3, used in the experiment. The left number in the leaves corresponds to the payoff for the player, the right number in the leaves corresponds to the payoff for the opponent.

2.4 Forward reasoning plus backtracking

Forward reasoning can be a fast strategy if the highest value is in the top part of the tree and there is no need for higher-order reasoning. However, all experimental trials contained the need for higher-order reasoning to find the best solution. Meijering et al. (2012) suggested that the participants were using forward reasoning plus backtracking. Forward reasoning plus backtracking is a combination of using forward reasoning to find the highest payoff and then using backtracking to determine if this payoff is attainable, see section 1.2. Algorithm 1 shows our implementation of the forward reasoning plus backtracking strategy on the Marble Drop Game. The algorithm is based on the description of forward reasoning plus backtracking by [Meijering et al. (2012)], but it is more generic than the algorithm used by [Meijering et al. (2012)]. Our algorithm can calculate the number of forward reasoning plus backtracking steps for any binary 2-player game tree.

We calculated the number of steps participants should make when they were using forward reasoning plus backtracking for every trial of the experimental data. We count a step as attention to a value. For example, comparing

Algorithm 1 Determine forward reasoning plus backtracking steps. Where m is the number of bins, P_n is the payoff for the player at bin n, O_n is the payoff for the opponent at bin n. This algorithm is more generic

Require: $P_n \in \{1:m\}$ **and** $O_n \in \{1:m\}$
Ensure: all P_n are unique **and** all O_n are unique
$n \Leftarrow 1$ {start with forward reasoning at the first bin}
$Steps \Leftarrow 1$
while not max P_n **do**
  $n \Leftarrow n + 1$ **and** $Steps \Leftarrow Steps + 1$ {While the highest payoff is not found continue with the next bin}
  if max P_n **and** max O_n **then**
    $Steps \Leftarrow Steps + 1$ {if the highest payoff of both players in in this bin there is no need for backtracking}
    return $Steps$
  end if
end while
$High \Leftarrow n$ {Remember the highest bin for the player}
$Back \Leftarrow m$
$n \Leftarrow m - 1$ {Start backtracking at the last two nodes}
while $Back \neq High$ **and** $n > 0$ **do**
  if trapdoor(n) = player **then**
    if $P_{Back} > P_n$ **then**
      $Back \Leftarrow Back$ {Node $Back$ has the highest score for the player, therefore the nodes can be substituted by node $Back$}
    else if $P_{Back} < P_n$ **then**
      $Back \Leftarrow n$ {Node n has the highest payoff for the player, therefore the nodes can be substituted by node n}
    end if
  else if trapdoor(n) = opponent **then**
    if $O_{Back} > O_n$ **then**
      $Back \Leftarrow Back$ {Node $Back$ has the highest score for the opponent, therefore the nodes can be substituted by node $Back$}
    else if $O_{Back} < O_n$ **then**
      $Back \Leftarrow n$ {Node n has the highest payoff for the opponent, therefore the nodes can be substituted by node n}
    end if
  end if
  $n \Leftarrow n - 1$
  $Steps \Leftarrow Steps + 2$ {There are two payoffs compares, hence this takes 2 steps}
end while
return $Steps$ {Return the number of Steps for forward reasoning plus backtracking}

two values in bins of the Marble Drop Game would be two steps, since both values need to be attended in order to compare them. Using this method, the trials were divided into three classes: 5, 6 and 8 steps, accordingly.

Figure 2 shows the payoff structures for type 1 and type 3 of the Marble Drop Game. As an example, we will walk through both types using forward reasoning plus backtracking :

type 1:
At first, the player will attend all leaves until he finds his highest payoff. The highest payoff is in the fourth leaf, hence it takes 4 steps. After finding his highest payoff in the right leaf of node U, the player needs to compare the payoff for the opponent in this leaf with the payoff for the opponent in the left leaf of node T. There are 2 nodes to compare, hence it takes 2 steps. The highest payoff for the opponent is in the left leaf of node T, hence the opponent would never let the player reach his highest payoff. Therefore, the highest payoff is not attainable. The player's last step is to compare his payoff in the left leaf of node T with his payoff in the left leaf of node S. There are 2 nodes to compare, hence it takes 2 steps. The left leaf of node S has the highest possible payoff. This takes a total of 8 steps.

type 3:
At first, the player will attend all leaves until he finds his highest payoff. The highest payoff is in the second leaf, hence it takes 2 steps. After finding his highest payoff in the left leaf of node T, the player needs to backtrack to know if this leaf is attainable. He starts comparing the left leaf of node U with the right leaf of node U. The right leaf of node U contains the highest payoff for the player. There are 2 nodes compared, hence it takes 2 steps. Next, he compares the payoff for the opponent in the left leaf of node T with the payoff for the opponent in the right leaf of node U, the highest payoff for the opponent is the left leaf of node T. There are 2 nodes compared, hence it takes 2 steps. This means the rational opponent would choose the left leaf at node T, which contains player 1's highest payoff, hence the highest payoff is attainable and no further backtracking is needed. This takes a total of 6 steps.

2.5 Lambda trees

The lambda-degree of a game tree with payoff structure is a suggested technique to indicate the complexity [Szymanik et al. (2013)], see section 1.3. We calculated the lambda-degree of the minimal subtree for all trials of the Marble Drop Game according to Definition 1. The minimal subtrees where created using Definition 2 and Definition 3. We have calculated the lambda-degree for the minimal subtree of all trials and we divided them into two categories: Λ_3^1 and Λ_2^1, accordingly.

Figure 2 shows the payoff structures for type 1 and type 3 of the Marble Drop Game. As an example, we will determine the lambda-degree of both types Using Definition 1, Definition 2 and Definition 3:

type 1:
The highest payoff for the player is in the right leaf of U, so the minimal subtree containing the highest payoff has 3 alterations between players. This means the minimal subtree is of lambda-degree Λ_3^1.

type 3:
The highest payoff for the player is in the left leaf of T, so the minimal subtree containing the highest payoff has 2 alterations between players. This means the minimal subtree is of lambda-degree Λ_2^1.

2.6 Models

To create models for explaining the reaction times, we used linear mixed-effects (LME) models. LME models can account for random effects of participants and unequal numbers of observations [Baayen, Davidson, and Bates (2008)]. Traditional methods (e.g. an ANOVA) cannot cope with missing or unequal numbers of variables. In order to validate our hypotheses we created two models, one for the H_Λ hypothesis and one for the H_{FRB} hypothesis. We used the log-transformed reaction times as dependent variable in both models, in order to obtain a close-to-normal distribution of the reaction times. To validate the H_{FRB} hypothesis, we used the following factors:

- An interaction between the calculated number of forward reasoning plus backtracking steps and whether or not the subject answered with the correct (i.e. optimal) solution. This way we account for the reaction times of trials where the subject answered the wrong solution, since these trials have a different reaction time from correct trials [Falkenstein, Hohnsbein, Hoormann, and Blanke (1991)].
- The trial number, in order to account for learning effects.
- The block number, in order to account for learning effects.

To validate the H_Λ hypothesis, we used the following factors:

- An interaction between the calculated lambda-degree of the minimal subtree and whether or not the subject answered with the correct (i.e. optimal) solution. This way we account for the reaction times of trials where the subject answered the wrong solution.
- The trial number, in order to account for learning effects.
- The block number, in order to account for learning effects.

For both models we used an automated selection algorithm that evaluated all factors one by one. The selection algorithm returns a model with the combination of factors from the base model that has the lowest Akaike information criterion (AIC) of all models. The AIC is a score for a model that is calculated as a trade-off between its complexity and the fit of the model and therefore is a way to select the best model out of a group of models [Akaike (1974)]. The models returned by the selection algorithm will be evaluated by AIC and R^2. To calculate the R^2, we used a variant of the R^2 statistic for LME models [Nakagawa and Schielzeth (2013)] which calculates the Marginal R^2 and the Condition R^2 statistic. The Marginal R^2 describes the proportion of variance explained by the fixed factors alone. The Conditional R^2 describes the proportion of variance explained by both the fixed and random factors.

2.7 Summary

Table 1 shows the categories for every type used in the experiment.

Table 1: Number of steps when using forward reasoning plus backtracking (FRB), lambda-degree (Lambda) and attainability for every type in the Marble Drop Game

type	FRB steps	Lambda	Attainable
1	8	Λ_3^1	No
2	8	Λ_2^1	No
3	6	Λ_2^1	Yes
4	8	Λ_2^1	No
5	8	Λ_2^1	No
6	8	Λ_3^1	No
7	5	Λ_3^1	Yes
8	6	Λ_2^1	Yes
9	6	Λ_2^1	Yes
10	8	Λ_3^1	No
11	6	Λ_3^1	Yes
12	5	Λ_3^1	Yes
13	8	Λ_3^1	No
14	6	Λ_3^1	Yes
15	6	Λ_2^1	Yes
16	8	Λ_2^1	No

3 Results

Table 2 contains the AIC, Marginal R^2, Conditional R^2 and factors of three models. Model 1 and model 2 are the models we originally created for validating our hypotheses. Model 3 is the model with the lowest AIC of any possible combination of the factors of model 2. Model 1 already has the lowest AIC of any possible combination of the factors of this model. Table 3 contains the p-values and the factors of model 1. Table 4 contains the p-values and the factors of model 2. Table 5 contains the p-values and the factors of model 3. We can see that the Correct factor is left out in model 3; we can also see that the Correct factor had little to none influence in model 2.

When applying a Chi-squared test on model 1 and model 3, we find that model 1 has the best fit, with model 1 $\chi^2 = 23.426$ and $p < 8.187e^{-6}$. Model 3 has no significant outcome.

4 Discussion

4.1 H_Λ: Lambda hypotheses

Model 3 is the selected model for evaluating H_Λ. This model does not predict the reaction times as well as the best model. The conditional R^2 shows us it can only account for 38.49% of the reaction time values. We can see that there is an influence of the lambda-degree of the minimal subtree on the reaction time. However, if we look at Table 5, we see that the estimate for this factor is negative. This means the higher the lambda-degree of the minimal subtree, the lower the reaction time. This is the exact opposite of our hypothesis, so we cannot confirm our H_Λ hypothesis.

4.2 H_{FRB}: FRB hypotheses

Model 1 is the selected model for evaluating H_{FRB}. This model is the best predictor for the reaction times. The conditional R^2 shows us it can account for 39.84% of the reaction time values. If we look at Table 3 we can see the interaction between forward reasoning plus backtracking steps and correct has positive estimate. This tells us a higher number of steps combined with a correct answer means a higher reaction time. Trial has a negative estimate, this means that subjects were faster at the end of the experiment than at the beginning. This is the learning effect we accounted for. The Block effect is positive, this means that subjects were slower in the second block than in the first block. The interaction between forward reasoning plus backtracking steps and correct has a p-value of $p < 0.00974$, this means we can confirm our H_{FRB} hypothesis. The number of forward reasoning plus backtracking steps is a good predictor for the reaction times of subjects, thus suggesting subjects use the forward reasoning plus backtracking strategy.

4.3 Overview

The research done in this study yielded insight in the strategy humans use when solving turn-based games with a binary payoff structure. Based on earlier research [Meijering et al.

Table 2: Calculated AIC, Marginal R^2 and Conditional R^2 values of all models. The factors of the models are a combination of: forward reasoning plus backtracking steps (FRB), lambda-degree (Lambda), the trial number (Trial), the block number (Block) and whether or not the participant gave the correct (i.e. optimal) answer (Correct). In all models, we use subject as a random effect (1|Subj).

Model	Factors	Marginal R^2	Conditional R^2	AIC
1	$FRB * Correct + Trial + Block + (1\|Subj)$	6.25%	39.48%	2827.793
2	$Lambda * Correct + Trial + Block + (1\|Subj)$	5.08%	38.49%	2850.049
3	$Lambda + Trial + Block + (1\|Subj)$	5.02%	38.49%	2847.219

Table 3: Fixed effect of model 1

	Estimate	St. Error	df	t value	Pr(> |t|)	
(Intercept)	9.15E+000	5.33E-001	1.08E+003	17.188	<2E-016	***
FRB	-6.14E-002	7.64E-002	1.20E+003	-0.803	0.42207	
Correct	-1.55E+000	5.30E-001	1.20E+003	-2.932	0.00343	**
Trial	-1.35E-002	2.33E-003	1.20E+003	-5.819	7.61E-009	***
Block	2.46E-001	4.31E-002	1.20E+003	5.719	1.35E-008	***
FRB:Correct	2.04E-001	7.89E-002	1.20E+003	2.589	0.00974	**

Table 4: Fixed effect of model 2

	Estimate	St. Error	df	t value	Pr(> |t|)	
(Intercept)	9.60E+000	4.40E-001	9.12E+002	21.802	<2E-016	***
Lambda	-3.92E-001	1.86E-001	1.20E+003	-2.113	0.0348	*
Correct	-3.88E-001	4.34E-001	1.20E+003	-0.895	0.3712	
Trial	-1.31E-002	2.35E-003	1.20E+003	-5.593	2.76E-008	***
Block	2.40E-001	4.35E-002	1.20E+003	5.515	4.28E-008	***
Lambda:Correct	1.46E-001	1.92E-001	1.20E+003	0.764	0.4449	

Table 5: Fixed effect of model 3

	Estimate	St. Error	df	t value	Pr(> |t|)	
(Intercept)	9.26E+000	1.85E-001	7.70E+001	50.16	<2E-016	***
Lambda	-2.61E-001	4.33E-002	1.20E+003	-6.031	2.16E-009	***
Block	2.37E-001	4.33E-002	1.20E+003	5.462	5.71E-008	***
Trial	-1.32E-002	2.35E-003	1.20E+003	-5.627	2.29E-008	***

(2012)][Szymanik et al. (2013)] we made models of different strategies subjects could use to best explain the reaction times and thus the strategy. We created a model that predicts the difficulty of a turn-based two player game with a binary payoff structure for a player using the forward reasoning plus backtracking strategy. [Meijering et al. (2012)] also found that participants were using forward reasoning plus backtracking when studying eye-tracking data from subjects playing the Marble Drop Game. By creating this model, we made it theoretically plausible that the forward reasoning plus backtracking strategy is indeed the strategy people use when solving turn-based games with a binary payoff structure.

4.4 Outlook

Other explanations that better suit our results are possible. We mostly focused on forward reasoning plus backtracking. However, there might be other strategies people could use that better suit our results. As a suggestion for further research we would recommend the study of other possible strategies for solving turn-based games. Another suggestion is the study of the forward reasoning plus backtracking hypothesis using other experimental methods than the Marble Drop Game. It would be interesting to study of turn-based games with larger binary payoff structures (e.g. six leaves instead of the four leaves we used). It would also be interesting to study if turn-based games with non-binary payoff structures yield the same result.

References

Hirotugu Akaike. A new look at the statistical model identification. *Automatic Control, IEEE Transactions on*, 19(6):716–723, 1974.

Ian Apperly. *Mindreaders: The Cognitive Basis of "Theory of Mind"*. Psychology Press, 2010.

R Harald Baayen, Douglas J Davidson, and Douglas M Bates. Mixed-effects modeling with crossed random effects for subjects and items. *Journal of Memory and Language*, 59 (4):390–412, 2008.

Jelmer P Borst, Niels A Taatgen, and Hedderik van Rijn. The problem state: A cognitive bottleneck in multitasking. *Journal of Experimental Psychology: Learning, memory, and cognition*, 36(2):363, 2010.

M Falkenstein, J Hohnsbein, J Hoormann, and L Blanke. Effects of crossmodal divided attention on late ERP components. II. error processing in choice reaction tasks. *Electroencephalography and clinical neurophysiology*, 78(6):447–455, 1991.

Liesbeth Flobbe, Rineke Verbrugge, Petra Hendriks, and Irene Krämer. Children's application of theory of mind in reasoning and language. *Journal of Logic, Language and Information*, 17(4):417–442, 2008.

Helen L Gallagher and Christopher D Frith. Functional imaging of 'theory of mind. *Trends in Cognitive Sciences*, 7(2):77–83, 2003.

Sujata Ghosh and Ben Meijering. On combining cognitive and formal modeling: A case study involving strategic reasoning. In *Proceedings of the Workshop on Reasoning About Other Minds (RAOM 2011). CEUR Workshop Proceedings*, volume 751, pages 79–92, 2011.

Wayne D Gray, Chris R Sims, Wai-Tat Fu, and Michael J Schoelles. The soft constraints hypothesis: A rational analysis approach to resource allocation for interactive behavior. *Psychological review*, 113(3):461, 2006.

Aviad Heifetz. *Game Theory: Interactive Strategies in Economics and Management.* Cambridge University Press, 2012.

Richard D McKelvey and Thomas R Palfrey. An experimental study of the centipede game. *Econometrica: Journal of the Econometric Society*, pages 803–836, 1992.

Ben Meijering, Hedderik Van Rijn, Niels Taatgen, and Rineke Verbrugge. I do know what you think I think: Second-order theory of mind in strategic games is not that difficult. In *Proceedings of the 33rd Annual Conference of the Cognitive Science Society, Cognitive Science Society, Austin (TX)*, pages 2486–2491, 2011.

Ben Meijering, Hedderik van Rijn, Niels A Taatgen, and Rineke Verbrugge. What eye movements can tell about theory of mind in a strategic game. *PloS ONE*, 7(9):e45961, 2012.

Shinichi Nakagawa and Holger Schielzeth. A general and simple method for obtaining R2 from generalized linear mixed-effects models. *Methods in Ecology and Evolution*, 4(2):133–142, 2013.

Josef Perner and Heinz Wimmer. "John thinks that Mary thinks that" attribution of second-order beliefs by 5-to 10-year-old children. *Journal of experimental child psychology*, 39 (3):437–471, 1985.

Jakub Szymanik, Ben Meijering, and Rineke Verbrugge. Using intrinsic complexity of turn-taking games to predict participants' reaction times. In *Proceedings of the 35th Annual Conference of the Cognitive Science Society, Cognitive Science Society, Austin (TX)*, 2013.

Rineke Verbrugge. Logic and social cognition: The facts matter and so do computational models. *Journal of Philosophical Logic*, 38 (6):649–680, 2009.

Attachment 1: Pay-off trees of all types of the Marble Drop Game:
The left number in the leaves corresponds to the payoff for the participant (player 1), the right number in the leaves corresponds to the payoff for the opponent (player 2).

Type 1:

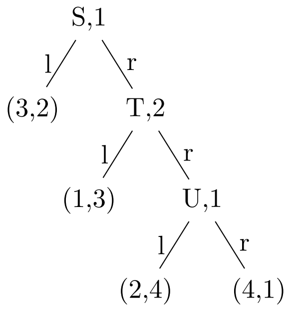

Type 2:

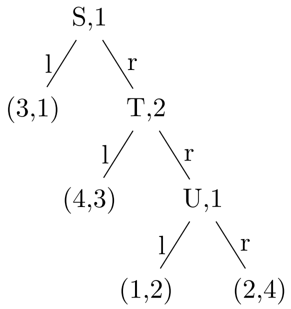

Type 3:

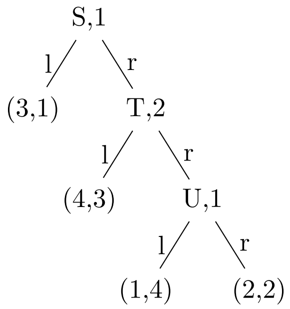

Type 4:

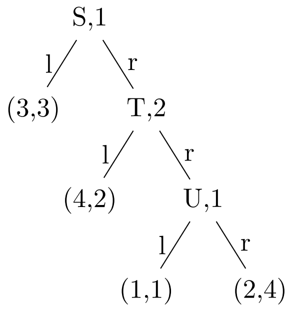

Type 5:

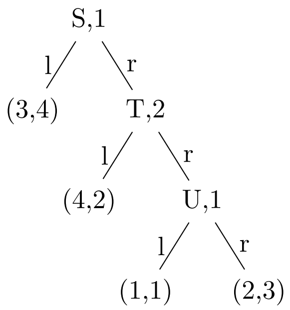

Type 6:

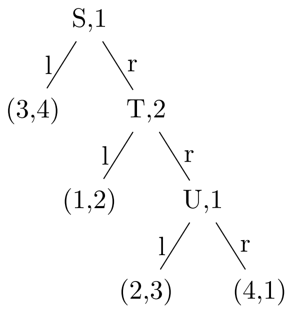

Type 7:

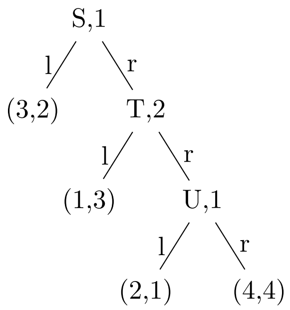

Type 8:

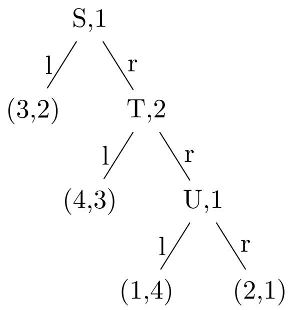

Type 9:

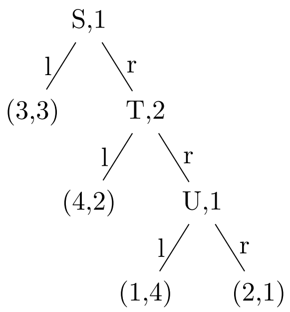

Type 10:

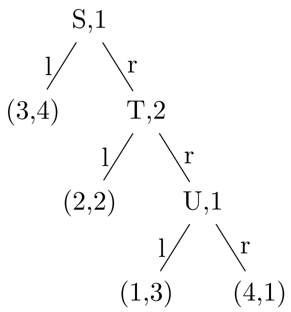

Type 11:

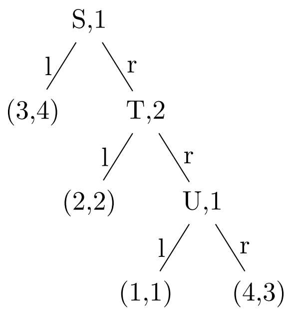

Type 12:

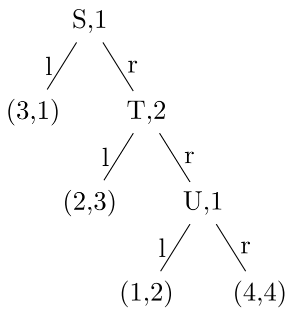

Type 13:

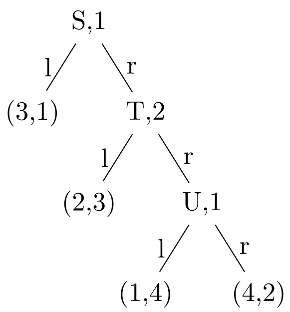

Type 14:

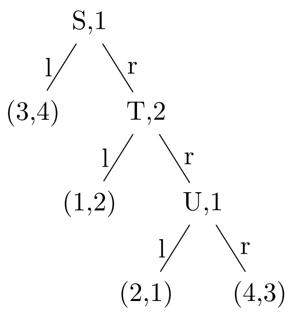

Type 15:

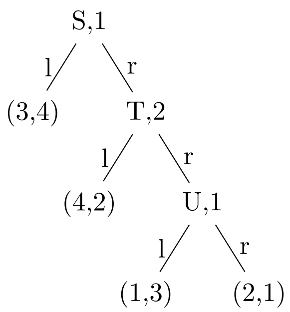

Type 16: